Parameterized Extension Complexity of Independent Set and Related Problems

Jakub Gajarský[a], Petr Hliněný[a], Hans Raj Tiwary[b]

[a]*Faculty of Informatics, Masaryk University Brno, Czech Republic.*
[b]*KAM/ITI, MFF, Charles University in Prague, Czech Republic.*

arXiv:1511.08841v2 [cs.CC] 18 Feb 2016

Abstract

Let G be a graph on n vertices and $\mathrm{STAB}_k(G)$ be the convex hull of characteristic vectors of its independent sets of size at most k. It is known that optimizing over $\mathrm{STAB}_k(G)$ is $W[1]$-hard and is FPT tractable for graphs of bounded expansion. We show analogous results for the extension complexity of $\mathrm{STAB}_k(G)$. In particular, we show that when G is a graph from a class of bounded expansion then $\mathrm{xc}(\mathrm{STAB}_k(G)) \leqslant \mathcal{O}(f(k) \cdot n)$ for some function f (depending only on the class). This result can be extended in a simple way to a wide range of similarly defined graph polytopes. In case of general graphs we show that there is *no function* f such that, for all natural numbers k and for all graphs on n vertices, the extension complexity of $\mathrm{STAB}_k(G)$ is at most $f(k) \cdot n^{\mathcal{O}(1)}$.

Keywords: FPT extension complexity, independent set, FO logic

1. Introduction

Polyhedral (aka LP) formulations of combinatorial problems belong to the basic toolbox of combinatorial optimization. In a nutshell, a set of feasible solutions of some problem is suitably encoded by a set of vectors, whose convex hull forms a polytope over which one can then optimize using established tools. A polytope Q is said to be an *extended formulation* or *extension* of a polytope P if P is a projection of Q. Measuring the size of a polytope by the minimum number of inequalities required to describe it, one can define the extension complexity of a polytope to be the smallest size extension of the polytope. This notion has a rich history in combinatorial optimization where by adding extra variables one can sometimes obtain significantly smaller polytopes. For some recent survey on extended formulations in the context of combinatorial optimization and integer programming see [1, 2, 3, 4].

Since linear (or indeed convex) optimization of a polytope P can instead be indirectly done by optimizing over an extended formulation of P, this concept provides a powerful model for solving many combinatorial problems. Various Linear Program (LP) solvers exist today that perform quite well in practice and it is desirable if a problem can be modeled as a small-sized polytope over which one can use an existing LP solver for linear optimization. However, in recent year super-polynomial lower bound on the extension complexity of polytopes associated with many combinatorial problems have been established. These bounds have been generalized to various settings, such as convex extended formulations, approximation algorithms, etc. These results are too numerous for a comprehensive listing but we refer the interested readers to some of the landmark papers in this regard [5, 6, 7, 8].

Many of the recent lower bounds on the extension complexity of various combinatorial polytopes mimic the computational complexity of the underlying problem. For example, it is known that the *extension complexities of polytopes* related to various NP-hard problems are super-polynomial [5, 9, 10, 6]. One satisfying feature of these lower bounds is that they are independent of traditional complexity-theoretic assumptions such as $P \neq NP$. Though, there also exist polytopes corresponding to polynomial

Email addresses: xgajar@fi.muni.cz (Jakub Gajarský), hlineny@fi.muni.cz (Petr Hliněný), hansraj@kam.mff.cuni.cz (Hans Raj Tiwary)

time solvable optimization problems whose extension complexity is super-polynomial. In particular, the perfect matching polytope was shown to have super-polynomial extension complexity by Rothvoß [7].

Notwithstanding the latter weakness, one can ask the related questions in the realm of *parameterized complexity theory*. In this rapidly grown field each problem instance comes additionally equipped with an integer parameter, and the "efficient" class denoted by FPT (*fixed-parameter tractable*) is the one of problems solvable, for every fixed value of the parameter, in polynomial time of degree independent of the parameter. See Section 2 for details.

Similarly as parameterized complexity provides a finer resolution of algorithmic tractability of problems, parameterized polytopes extensions can provide a finer resolution of extension complexities of polytopes of the problems. We follow this direction of research with a case study of the *independent-set polytope* of a graph, naturally parameterized by the solution size. We confirm that the extension complexity of the independent-set polytope indeed mimics the parameterized complexity of the underlying problem, in the following sense:

- We prove that this polytope does not have FPT extension for all graphs (Section 3), but
- an FPT extensions of it exists on every graph class of bounded expansion (Section 4).

We conclude the paper with a straightforward extension of this result to many other similarly defined problems, namely those expressible in existential FO logic of graphs (Section 5), followed by some further thoughts on the topic (Section 6).

2. Preliminaries

We follow standard terminology of graph theory and consider finite simple undirected graphs. We refer to the vertex and edge sets of a graph G as to $V(G)$ and $E(G)$, respectively. In particular, an *independent set* X of vertices of a graph is such that no two elements of X are joined by an edge. By a *cut* in a graph G we mean an edge cut, that is a minimal (by inclusion) set of edges $C \subseteq E(G)$ such that $G \setminus C$ has more connected components than G.

For fundamental concepts of parameterized complexity we refer the readers, e.g., to the monograph [11]. Here we just very briefly recall the needed notions. Considering a problem $\mathcal{P}$ with input of the form $(x, k) \in \Sigma^* \times \mathbb{N}$ (where k is a *parameter*), we say that $\mathcal{A}$ is *fixed-parameter tractable* (shortly FPT) if there is an algorithm solving $\mathcal{A}$ in time $f(k) \cdot n^{\mathcal{O}(1)}$ where f is an arbitrary computable function. In the (parameterized) *k-independent set problem* the input is (G, k) where G is a graph and $k \in \mathbb{N}$, and the question is whether G has an independent set of size at least k.

There is no known FPT algorithm for the k-independent set problem in general and, in fact, the theory of parameterized complexity [11] defines complexity classes $W[t]$, $t \geq 1$, such that the k-independent set problem is complete for $W[1]$. Problems that are $W[1]$-hard do not admit an FPT algorithm unless the Exponential Time Hypothesis fails.

2.1. Fixed-parameter extension complexity

The *size* of a polytope P, denoted by $\mathrm{size}(P)$, is defined to be the number of facets of P, which is the minimum number of inequalities needed to describe P if it is full-dimensional. A polytope Q is called an *extension* of a polytope P if P can be obtained as a linear projection of Q. As a shorthand we will say that in this case Q is an EF of P. As noted in the Introduction, the following is a useful notion:

Definition 2.1 (Extension complexity)**.** The *extension complexity of a polytope P*, denoted by $\mathrm{xc}(P)$, is defined to be the size of the smallest extension. More precisely,

$$\mathrm{xc}(P) := \min_{Q \text{ an EF of } P} \mathrm{size}(Q).$$

In the context of *fixed-parameter extension complexity*, we deal with families of polytopes $\mathcal{P}_n$ where $n \in \mathbb{N}$, and a parameter k. For example, for the independent set problem parameterized by a nonnegative integer k, the family $\mathcal{P}_n$ could be the family of k-independent set polytopes (cf. Subsection 2.2) for a family of n-vertex graphs. The prime question is whether there exists a computable function f such that $\mathrm{xc}(P) \leq f(k) \cdot n^{\mathcal{O}(1)}$ for all k, n and all $P \in \mathcal{P}_n$. As a shorthand we will say in such case that the collection of families $\{\mathcal{P}_n : n \in \mathbb{N}\}$ has *FPT extension complexity* (analogously to the FPT complexity class).

Buchanan in a recent article [12] studied the fixed-parameter extension complexity of the k-vertex cover problem, and proved that for any graph G with n vertices, the k-vertex cover polytope has extension complexity at most $\mathcal{O}(c^k n)$ for some constant $c < 2$. Hence this is a nontrivial example of a polytope class with FPT extension complexity. Buchanan also raised the question whether the k-independent set polytope (Definition 2.4) admits an FPT extension. We answer this in the negative in Theorem 3.6. Note that our negative answer does *not rely* on any complexity theoretical assumptions (such as $FPT \neq W[1]$).

We also look at the positive side of the k-independent set problem. It is known that this problem admits an FPT algorithm (w.r.t. k) on quite rich restricted graph classes, e.g., on classes of bounded expansion [13] (see Subsection 2.3 for the definition). While this finding, in general, does not imply anything about the extension complexity of the k-independent set polytope, we manage to apply the tools of [13] in our setting, and confirm – in Theorem 4.3 – FPT extension complexity of the k-independent set polytope on any graph class of bounded expansion. We also study a meta-generalization of this result in Section 5.

In the course of proving aforementioned Theorems 3.6 and 4.3, we are going to use the following established results on the topic of polytope extension complexity.

Theorem 2.2 (Balas [14])**.** *Let $P_1, P_2, \ldots, P_s$ be polytopes and let $P := \mathrm{conv}(\bigcup_{i=1}^{s} P_i)$. Then,* $\mathrm{xc}(P) \leqslant s + \sum_{i=1}^{s} \mathrm{xc}(P_i)$.

For a graph G, a cut vector is a 0/1 vector of length $|E(G)|$ whose coordinates correspond to whether an edge of G is in a cut $C \subseteq E(G)$ or not. The *cut polytope* is then a convex hull of all the cut vectors of G. Our negative result relies on the following lower bound.

Theorem 2.3 (Fiorini et al. [5])**.** *The extension complexity of the cut polytope of the complete graph K_n on n vertices is $2^{\Omega(n)}$.*

2.2. The k-independent set polytope

Let G be a graph on n vertices. Every subset of vertices of G can be encoded as a characteristic vector of length n. That is, for a subset $S \subseteq V$, define the characteristic vector χ^S as follows:

$$\chi^S_v = \begin{cases} 1 & \text{if } v \in S \\ 0 & \text{otherwise} \end{cases}$$

Definition 2.4 (independent set polytope)**.** The *k-independent set polytope of G*, denoted by $\mathrm{STAB}_k(G)$, is defined to be the convex hull of the characteristic vectors of every independent set of size at most k. That is,

$$\mathrm{STAB}_k(G) = \mathrm{conv}\left(\left\{\chi^S \in \{0,1\}^n |\ S \subseteq V \text{ is an independent set of } G; |S| \leqslant k\right\}\right).$$

In case that $k = n$ we simply speak about the *independent set polytope of G*; $\mathrm{STAB}(G)$.

Alternatively, one could define the polytope $\mathrm{STAB}^{=}_k(G)$ to be the convex hull of all independent sets of size exactly equal to k. That is,

$$\mathrm{STAB}^{=}_k(G) = \mathrm{conv}\left(\left\{\chi^S \in \{0,1\}^n |\ S \subseteq V \text{ is an independent set of } G; |S| = k\right\}\right).$$

To simplify our situation, we note the following:

Lemma 2.5. $\mathrm{xc}(\mathrm{STAB}_k^=(G)) \leqslant \mathrm{xc}(\mathrm{STAB}_k(G)) \leqslant k + \sum_{i=0}^{k} \mathrm{xc}(\mathrm{STAB}_i^=(G))$.

Proof. Clearly, $\mathrm{STAB}_k^=(G)$ is a face of $\mathrm{STAB}_k(G)$. Therefore, $\mathrm{xc}(\mathrm{STAB}_k^=(G)) \leqslant \mathrm{xc}(\mathrm{STAB}_k(G))$.

On the other hand, $\mathrm{STAB}_k(G) = \mathrm{conv}(\bigcup_{i=1}^{k} \mathrm{STAB}_i^=(G))$, and therefore $\mathrm{xc}(\mathrm{STAB}_k(G)) \leqslant k + \sum_{i=0}^{k} \mathrm{xc}(\mathrm{STAB}_i^=(G))$ by Theorem 2.2. □

The above Lemma 2.5 shows that any bounds (lower or upper) that are valid for $\mathrm{xc}(\mathrm{STAB}_k^=(G))$ are also asymptotically valid for $\mathrm{xc}(\mathrm{STAB}_k(G))$. Therefore, in the rest we will use the notation $\mathrm{STAB}_k(G)$ to actually mean $\mathrm{STAB}_k^=(G)$, to keep the notation uncluttered. We stress that this is just for simplicity of notation and does not cause any loss of generality.

We shall also use the following result.

Theorem 2.6 (Buchanan and Butenko [15])**.** *The extension complexity of a graph's independent set polytope is $\mathcal{O}(2^{tw}n)$, where n is the number of vertices and tw denotes its treewidth.*

Note that Buchanan and Butenko give an explicit description of an extension of the independent set polytope.

2.3. Sparsity and bounded expansion

A useful toolbox in our research is the theory of sparse graph classes, largely developed by Nešetřil and Ossona de Mendez. We follow their monograph [16].

We start by defining the notion of edge contraction. Given an edge $e = uv$ of a graph G, we let G/e denote the graph obtained from G by *contracting* the edge e, which amounts to deleting the endpoints of e, introducing a new vertex w_e and making it adjacent to all vertices in the union of the neighborhoods of u and v (excluding u, v themselves). A *minor* of G is a graph obtained from a subgraph of G by contracting zero or more edges. In a more general view, if H is isomorphic to a minor of G, then we call H a minor of G as well, and we write $H \preceq G$.

Alternatively, H is a minor of G if there exists a bijection $\psi\colon V(H) \to \{V_1, \ldots, V_p\}$ where $V_1, \ldots, V_p$ are pairwise disjoint subsets of $V(G)$ inducing connected subgraphs of G, and $uv \in E(H)$ only if there is an edge in G with an endpoint in each of $\psi(u)$ and $\psi(v)$. If, moreover, it is required that each subgraph $G[V_i]$ has radius at most d, meaning that there exist $c_i \in V_i$ (a center) such that every vertex in V_i is within distance at most d from c_i in $G[V_i]$; then H is called a *shallow minor at depth d* of G (shortly, a *d-shallow minor*).

Note that if $u, v \in V(H)$ in a d-shallow minor, and $u_1 \in \psi(u)$ and $v_1 \in \psi(v)$, then $d_G(u_1, v_1) \leq (2d+1) \cdot d_H(u, v)$. The class of d-shallow minors of G is denoted by $G\nabla d$, and this is extended to graph classes $\mathcal{G}$ as well; $\mathcal{G}\nabla d = \bigcup_{G \in \mathcal{G}} G\nabla d$.

One of the most prominent [16] notions of "sparsity" for graph classes is the following one:

Definition 2.7 (Grad and bounded expansion [17])**.** Let $\mathcal{G}$ be a graph class. Then the *greatest reduced average density* of $\mathcal{G}$ with *rank* d is defined as

$$\nabla_d(\mathcal{G}) = \sup_{H \in \mathcal{G}\nabla d} \frac{|E(H)|}{|V(H)|}.$$

A graph class $\mathcal{G}$ has *bounded expansion* if there exists a function $f : \mathbb{N} \to \mathbb{R}$ (called the *expansion function*) such that for all $d \in \mathbb{N}$, $\nabla_d(\mathcal{G}) \leq f(d)$.

We provide a brief informal explanation of Definition 2.7. A graph to be considered "sparse" should not, in particular, contain subgraphs with relatively many edges. Since $G\nabla 0$ is the set of all subgraphs of G, this is captured by $2\nabla_0(G)$ being the maximum average degree over all subgraphs of G. However, the definition requires more; even after contracting edges up to limited depth d, the resulting shallow minors stay free of relatively dense subgraphs, with the maximum average degree bounded by $2\nabla_d(\mathcal{G}) \leq 2f(d)$.

For example, the class $\mathcal{P}$ of all planar graphs has bounded expansion (even with a constant expansion function). On the other hand, a class $\mathcal{Q}$ obtained from all cliques by subdividing each edge twice, although also having relatively few edges, does not have bounded expansion since $\mathcal{Q}\nabla 1$ contains all graphs.

3. Lower Bound: Paired Local-Cut Graphs

We use a shorthand notation $[n] = \{1, 2, \ldots, n\}$. Given positive integers k and n, we define a graph called a *Paired Local-Cut Graph* and denoted $\mathrm{PLC}(k, n)$.

First we create $k2^{\lfloor \log n \rfloor}$ vertices labeled with tuples (i, S) for $i \in [k]$ and $S \subseteq [\lfloor \log n \rfloor]$. These vertices will be called *cut vertices*. Then we create $2\binom{k}{2}2^{2\lfloor \log n \rfloor}$ vertices labeled with tuples (i, j, S_1, S_2) where $1 \leqslant i \neq j \leqslant k$ and $S_1, S_1 \subseteq [\lfloor \log n \rfloor]$. These vertices will be called *pairing vertices*.

We add edges to these vertices of $\mathrm{PLC}(k, n)$ as follows. For each fixed $i \in [k]$ we add the edges between all cut nodes that have labels (i, S). Furthermore, for each fixed pair $i, j \in [k]$ we add the edges between all pairing nodes that have labels (i, j, S_1, S_2). Finally, let u be a cut vertex labeled (i, S) and let v be a pairing vertex labeled (j_1, j_2, S_1, S_2). If $i = j_1$ but $S \neq S_1$ we add edge uv. Symmetrically, if $i = j_2$ but $S \neq S_2$ we add edge uv.

For ease of exposition we will identify vertices of $\mathrm{PLC}(k, n)$ with their labels whenever convenient.

Observation 3.1. The number of vertices of the graph $\mathrm{PLC}(k, n)$ equals $2\binom{k}{2}2^{2\lfloor \log n \rfloor} + k2^{\lfloor \log n \rfloor} \leqslant (kn)^2$.

Observation 3.2. Let (i, S) and (j_1, j_2, S_1, S_2) be two vertices of $\mathrm{PLC}(k, n)$ that are not joined by an edge. If $i = j_1$ then $S = S_1$, and if $i = j_2$ then $S = S_2$.

This together with the next lemma will ensure that in any independent set I of $\mathrm{PLC}(k, n)$ that has size k^2, every index $i \in [k]$ can be uniquely associated with a subset $S_i \subseteq \lfloor \log n \rfloor$.

Lemma 3.3. *Let I be an independent set in $\mathrm{PLC}(k, n)$. Then, $|I| \leqslant k^2$. Moreover, an equality holds if and only if I contains exactly one cut vertex for each $1 \leqslant i \leqslant k$ and exactly one pairing vertex for each $1 \leqslant i \neq j \leqslant k$.*

Proof. Clearly, the set I can contain at most k cut vertices – at most one vertex (i, S_i) for each $1 \leqslant i \leqslant k$. Also, I can contain at most $2\binom{k}{2} = k^2 - k$ pairing vertices – at most one vertex (i, j, S_i, S_j) for each ordered pair $1 \leqslant i, j \leqslant k$. □

We will now relate the vertices of $\mathrm{STAB}_{k^2}(\mathrm{PLC}(k, n))$ with the vertices of the *polytope* $\mathrm{CUT}(K_r)$ where $r = k \lfloor \log n \rfloor$, to be defined as follows. We denote the vertices and edges of the complete graph K_r (on r vertices) by V_r and E_r, respectively. We group the vertices of K_r into k groups, each of size $\lfloor \log n \rfloor$, and label the vertices v^i_j where $1 \leqslant i \leqslant k$ and $1 \leqslant j \leqslant \lfloor \log n \rfloor$. Finally, we order the vertices lexicographically according to their labels.

A cut vector of K_r – corresponding to a cut C – is a 0/1 vector of length $\binom{r}{2}$ whose coordinates correspond to whether an edge of K_r is in the cut C or not. The edges of K_r are labeled with pairs (i_1, j_1, i_2, j_2) where $1 \leqslant i_1, i_2 \leqslant k$; $1 \leqslant j_1, j_2 \leqslant \lfloor \log n \rfloor$, and $(i_1, j_1) \leqslant (i_2, j_2)$ lexicographically. So, if z is a cut vector corresponding to a given cut $C \subset E_r$, then $z_{i_1,j_1,i_2,j_2} = 1$ if and only if the edge (i_1, j_1, i_2, j_2) is in C. $\mathrm{CUT}(K_r)$ is the convex hull of all such cut vectors.

Similarly, an independent-set vector of $\mathrm{PLC}(k, n)$ – corresponding to an independent set I – is a 0/1 vector of length $2\binom{k}{2}2^{2\lfloor \log n \rfloor} + k2^{\lfloor \log n \rfloor}$ (see Observation 3.1) whose coordinates correspond to whether the corresponding vertex is in I or not. Recall that the cut vertices of $\mathrm{PLC}(k, n)$ are labeled with a pair consisting of an index from $[k]$, and a subset of $[\lfloor \log n \rfloor]$. Also, the pairing vertices of $\mathrm{PLC}(k, n)$ are labeled with a tuple consisting of two indices from $[k]$ and two subsets of $[\lfloor \log n \rfloor]$.

Let $\mathcal{C}$ be the set of all cuts in K_r, and let $\mathcal{I}$ be the set of all independent sets of size k^2 in $\mathrm{PLC}(k, n)$. Any cut $C \in \mathcal{C}$ creates a bipartition $(S, \overline{S})$ of the vertices of K_r. Recall that the vertices of K_r have been split in k groups. The partition $(S, \overline{S})$ thus induces a partition $(S_i, \overline{S}_i)$ within each of these groups.

Lemma 3.4. *For every pair of natural numbers (k, n) and $r = k \lfloor \log n \rfloor$ it holds that $\mathrm{CUT}(K_r)$ is a projection of $\mathrm{STAB}_{k^2}(\mathrm{PLC}(k, n))$.*

Proof. Let s denote the length of the independent set vectors of $\mathrm{STAB}_{k^2}(\mathrm{PLC}(k, n))$. That is $s = 2\binom{k}{2}2^{2\lfloor \log n \rfloor} + k2^{\lfloor \log n \rfloor}$. We describe an affine map $\pi : \mathbb{R}^s \to \mathbb{R}^{\binom{r}{2}}$ such that for every vertex C of $\mathrm{CUT}(K_r)$ there exists a vertex of I of $\mathrm{STAB}_{k^2}(\mathrm{PLC}(k, n))$ such that $\pi(I) = C$. Moreover, for every

vertex I of $\mathrm{STAB}_{k^2}(\mathrm{PLC}(k,n))$ we show that $\pi(I)$ is a vertex of $\mathrm{CUT}(K_r)$. Since π is an affine map, this will complete the proof.

First we identify the coordinates of $\mathbb{R}^s$ with vertices of $\mathrm{PLC}(k,n)$. To make it easy to refer to this identification we label the coordinates with tuples (i,j,S_1,S_2) defined as follows. For a coordinate corresponding to a cut vertex (i,S) we label the coordinate with (i,i,S,S). For a coordinate corresponding to a pairing vertex (i,j,S_1,S_2) we label the coordinate with the same label. Similarly, we identify the coordinates of $\mathbb{R}^{\binom{r}{2}}$ with the pairs of vertices of K_r, i.e, the coordinate corresponding to an edge between two distinct vertices $v^i_{\ell_1}$ and $v^j_{\ell_2}$ is to be labeled with the integer tuple (i,ℓ_1,j,ℓ_2), assuming that $v^i_{\ell_1} \leqslant v^j_{\ell_2}$ (that is, $i \leqslant j$ and if $i=j$ then $\ell_1 < \ell_2$). Also note that $1 \leqslant i,j \leqslant k$ and $1 \leqslant \ell_1,\ell_2 \leqslant \lfloor \log n \rfloor$.

Given a vector $y \in \mathbb{R}^s$ we define $\pi(y) := z \in \mathbb{R}^{\binom{r}{2}}$ where

$$z_{i_1,\ell_1,i_2,\ell_2} = \begin{cases} \displaystyle\sum_{\substack{S_1 \subseteq [\lfloor \log n \rfloor] \\ S_2 \subseteq [\lfloor \log n \rfloor] \\ \ell_1 \notin S_1 \\ \ell_2 \in S_2}} y_{i_1,i_2,S_1,S_2} + \sum_{\substack{S_1 \subseteq [\lfloor \log n \rfloor] \\ S_2 \subseteq [\lfloor \log n \rfloor] \\ \ell_1 \in S_1 \\ \ell_2 \notin S_2}} y_{i_1,i_2,S_1,S_2} & \text{if } i_1 \neq i_2, \\ \displaystyle\sum_{\substack{S \subseteq [\lfloor \log n \rfloor] \\ \ell_1 \notin S \\ \ell_2 \in S}} y_{i_1,i_1,S,S} + \sum_{\substack{S \subseteq [\lfloor \log n \rfloor] \\ \ell_1 \in S \\ \ell_2 \notin S}} y_{i_1,i_1,S,S} & \text{if } i_1 = i_2\,. \end{cases}$$

Let $y \in \mathbb{R}^s$ be a vertex of $\mathrm{STAB}_{k^2}(\mathrm{PLC}(k,n))$. That is, y is the characteristic vector of an independent set $I \in \mathcal{I}$. Since I is of size k^2, for every $1 \leqslant i \leqslant k$ exactly one cut vertex (i,S_i) is picked in I, by Lemma 3.3. Furthermore, by Observation 3.2, for any pairing vertex (i_1,i_2,S,S') picked in I the sets S,S' are unique for given i_1,i_2; if $i=i_1$ then $S=S_i$, and if $i=i_2$ then $S'=S_i$. Consider the subsets $S(I),\overline{S}(I)$ of vertices of K_r defined as follows:

$$\begin{aligned} S(I) &= \left\{ v^i_j \middle|\, 1 \leqslant i \leqslant k; j \in S_i \right\} \\ \overline{S}(I) &= \left\{ v^i_j \middle|\, 1 \leqslant i \leqslant k; j \notin S_i \right\} \end{aligned}$$

It is not difficult to see that $\pi(y)$ is exactly the characteristic vector of the cut defined by $S(I),\overline{S}(I)$ because $z_{i_1,\ell_1,i_2,\ell_2} = 1$ if and only if $v^{i_1}_{\ell_1}$ and $v^{i_2}_{\ell_2}$ do not both lie in $S(I)$ or both in $\overline{S}(I)$.

On the other hand, if z is the characteristic vector of a cut defined by subsets $S,\overline{S}$ of vertices of K_r, then we can define $S_i = S \cap \left\{ v^i_j \middle|\, 1 \leqslant j \leqslant \lfloor \log n \rfloor \right\}$ so that $\left\{ (i,S_i) | 1 \leqslant i \leqslant k \right\} \cup \left\{ (i,j,S_i,S_j) | 1 \leqslant i < j \leqslant k \right\}$ is an independent set I of $\mathrm{PLC}(k,n)$ whose size is k^2 and whose characteristic vector projects to z under π.

Hence π defines a projection from $\mathrm{PLC}(k,n)$ to $\mathrm{CUT}(K_r)$. □

Corollary 3.5. *There exists a constant $c' > 0$ such that for $k,n \in \mathbb{N}$,*

$$\mathrm{xc}\left(\mathrm{STAB}_{k^2}(\mathrm{PLC}(k,n))\right) \geqslant n^{c'k}.$$

Proof. By Lemma 3.4, $\mathrm{STAB}_{k^2}(\mathrm{PLC}(k,n))$ is an extended formulation of $\mathrm{CUT}(K_r)$ with $r = k\lfloor \log n \rfloor$. So any extended formulation of $\mathrm{STAB}_{k^2}(\mathrm{PLC}(k,n))$ is also an extended formulation of $\mathrm{CUT}(K_r)$. By Theorem 2.3, $\mathrm{xc}\left(\mathrm{CUT}(K_r)\right) \geqslant 2^{\Omega(r)}$. Therefore, $\mathrm{xc}\left(\mathrm{STAB}_{k^2}(\mathrm{PLC}(k,n))\right) \geqslant \mathrm{xc}\left(\mathrm{CUT}(K_r)\right) \geqslant 2^{\Omega(r)} \geqslant n^{c'k}$ for some constant $c' > 0$. □

We can now easily finish with the main result of this section.

Theorem 3.6. *There is no function $f : \mathbb{N} \to \mathbb{R}$ such that $\mathrm{xc}(\mathrm{STAB}_k(G)) \leqslant f(k) \cdot n^{\mathcal{O}(1)}$ for all natural numbers k and all graphs G on n vertices.*

Proof. Suppose, on the contrary, that such a function f does exist. That is, there is a constant c such that for every pair of natural numbers (ℓ, m) and for all m-vertex graphs G it holds that $\mathrm{xc}(\mathrm{STAB}_\ell(G)) \leqslant f(\ell) \cdot m^c$.

Given a pair (k, n) of natural numbers consider the graph $\mathrm{PLC}(k, n)$. By Corollary 3.5, we have that $\mathrm{xc}\left(\mathrm{STAB}_{k^2}(\mathrm{PLC}(k, n))\right) \geqslant n^{c'k}$ for some constant $c' > 0$. On the other hand, from our assumption for $\ell = k^2$ and $m \leqslant (kn)^2$ we have that $\mathrm{xc}\left(\mathrm{STAB}_{k^2}(G)\right) \leqslant f(k^2) \cdot (kn)^{2c}$. Therefore, $n^{c'k} \leqslant f(k^2) \cdot (kn)^{2c}$ and so

$$c'k \log n \leqslant \log f(k^2) + 2c(\log k + \log n)$$

$$(c'k - 2c) \log n \leqslant \log f(k^2) + 2c \log k.$$

Clearly, the latter cannot hold true for a sufficiently large but fixed parameter k and arbitrary n, a contradiction. Hence no such function f exists. □

4. Upper Bound: Bounded Expansion Classes

While Theorem 3.6 asserts that FPT extensions are not possible for the k-independent set polytopes of all graphs, there is still a good chance to prove a positive result for restricted classes of graphs. An example of such restriction is, by a simple modification of Theorem 2.6, presented by graph classes of bounded tree-width; although, this is somehow too restrictive. We show that in the case of k being a fixed parameter, one can go much further.

The underlying idea of our approach can be informally explained as follows. Imagine we can "guess", in advance, a (short) list of well-structured subgraphs of our graph such that every possible independent set is fully contained in at least one of them. Then we can separately construct an independent set polytope for each one of the subgraphs, and make their union at the end (Theorem 2.2). This ambitious plan indeed turns out to be viable for graph classes of bounded expansion (Definition 2.7), and the key to the success is a combination of a powerful structural characterization of bounded expansion (Theorem 4.2) with the size bound k on the independent sets.

In order to state the desired structural characterization, we need the notion of treedepth. In this context, a *rooted forest* is a disjoint union of rooted trees. The height of a rooted forest is the maximum distance from one of the forest's roots to a vertex in the same tree. The *closure* $\mathrm{clos}(\mathcal{F})$ of a rooted forest $\mathcal{F}$ is the graph with the vertex set $\bigcup_{T \in \mathcal{F}} V(T)$ and the edge set $\{xy : x \text{ is an ancestor of } y \text{ in a tree of } \mathcal{F}\}$. The *treedepth* $\mathrm{td}(G)$ of a graph G is the minimum height plus one of a rooted forest $\mathcal{F}$ such that $G \subseteq \mathrm{clos}(\mathcal{F})$.

For simplicity, we skip here the rather complicated definition of treewidth (which is never directly used in our arguments, anyway). Instead, we just use the following fact which is easy to establish directly from the definitions:

Observation 4.1. For any G, the treewidth of G is at most the treedepth of G minus one.

The amazing connection between graph classes of bounded expansion and treedepth is captured by the notion of low treedepth coloring: For an integer $d \geq 1$, an assignment of colors to the vertices of a graph G is a *low tree-depth coloring of order d* if, for every $s = 1, 2, \ldots, d$, the union of any s color classes induces a subgraph of G of treedepth at most s.

In particular, every low tree-depth coloring of G is a proper coloring of G (but not the other way round), and the union of any two color classes induces a forest of stars. The following result is crucial:

Theorem 4.2 (Nešetřil and Ossona de Mendez [13, 17])**.** *If $\mathcal{G}$ is a class of graphs of bounded expansion, then there is a function $N_{\mathcal{G}} : \mathbb{N} \to \mathbb{N}$ (depending on the expansion function of $\mathcal{G}$) such that for any graph $G \in \mathcal{G}$ and k, there exists a low tree-depth coloring of order k of G using at most $N_{\mathcal{G}}(k)$ colors. This coloring can be found in linear time for fixed k.*

We are now ready to state and prove the main theorem of this section.

Theorem 4.3. *Let $\mathcal{G}$ be any graph class of bounded expansion. Then there exists a computable function $f : \mathbb{N} \to \mathbb{N}$, depending on the expansion function of $\mathcal{G}$, such that*

$$\mathrm{xc}\big(\mathrm{STAB}_k(G)\big) \leqslant f(k) \cdot n$$

holds for every integer n and every n-vertex graph $G \in \mathcal{G}$. Moreover, an explicit extension of $\mathrm{STAB}_k(G)$ of size at most $f(k) \cdot n$ can be found in linear time for fixed k and $\mathcal{G}$.

Proof. Since $\mathcal{G}$ is a graph class of bounded expansion, by Theorem 4.2 we can for any $G \in \mathcal{G}$ and given k find an assignment $c : V(G) \to [N_{\mathcal{G}}(k)]$ such that c is a low tree-depth coloring of order k. Let $\mathcal{J}_k := \binom{[N_{\mathcal{G}}(k)]}{k}$ denote the set of k-element subsets of $[N_{\mathcal{G}}(k)]$, and let a subgraph $G_J \subseteq G$ where $J \in \mathcal{J}_k$, be defined as the subgraph of G induced on $\bigcup_{j \in \mathcal{J}_k} c^{-1}(j)$ – the color classes indexed by J.

Note the following two immediate facts:

a) by the definition, each G_J, $J \in \mathcal{J}_k$, is of treedepth at most $|J| = k$;

b) for every set $X \subseteq V(G)$ (independent or not) of size $|X| \leq k$, there is $J \in \mathcal{J}_k$ such that $X \subseteq V(G_J)$.

Consequently,

$$\mathrm{STAB}_k(G) = \mathrm{conv}\left(\bigcup\nolimits_{J \in \mathcal{J}_k} \mathrm{STAB}_k(G_J)\right)$$

and it is sufficient to bound the extension complexity of each $\mathrm{STAB}_k(G_J)$.

By (a) and Observation 4.1, $tw(G_J) \leq k-1$ and Theorem 2.6 applies here: $\mathrm{xc}\big(\mathrm{STAB}_k(G_J)\big) \leqslant \mathcal{O}(2^k \cdot |G_J|) \leqslant \mathcal{O}(2^k \cdot n) \leqslant c' 2^k \cdot n$ for a suitable constant c'. Then, by Theorem 2.2, we have

$$\begin{aligned}\mathrm{xc}\big(\mathrm{STAB}_k(G)\big) &\leqslant |\mathcal{J}_k| + \sum_{J \in \mathcal{J}_k} \mathrm{xc}\big(\mathrm{STAB}_k(G_J)\big) \\ &\leqslant |\mathcal{J}_k| \cdot (1 + c' 2^k \cdot n) \leqslant \binom{N_{\mathcal{G}}(k)}{k} (1 + c' 2^k) \cdot n \leqslant f(k) \cdot n\,.\end{aligned}$$

Note that this extended formulation can be constructed in linear time for fixed k since the low tree-depth coloring in Theorem 4.2 can be found in linear time, the extended formulation in Theorem 2.6 is explicit, and the extended formulation of union of polytopes can be constructed in linear time from the extensions of the component polytopes [14]. □

5. Generalizing the Upper Bound

The tools used in Section 4 for graph classes of bounded expansion apply not only to the particular case of independent sets of size at most k, but to a much wider range of problems. This can be formulated as a meta-result about graph polytopes definable in the first-order logic.

The *first-order logic of graphs* (abbreviated as FO) applies the standard language of first-order logic to a graph G viewed as a relational structure with the domain $V(G)$ and the single binary (symmetric) relation $E(G)$. For example, the formula $\iota(x_1, \ldots, x_k) \equiv \bigwedge_{i \neq j} \big(\neg edge(x_i, x_j) \wedge x_i \neq x_j\big)$ asserts that $\{x_1, \ldots, x_k\}$ is an independent set of size exactly k. A slightly more involved example describes a vertex cover tuple as $\gamma(x_1, \ldots, x_k) \equiv \forall y,z\big(edge(y, z) \to \bigvee_{i=1}^{k} (y = x_i \vee z = x_i)\big)$.

To any FO formula $\phi(x_1, \ldots, x_k)$ one can assign a graph polytope, although this task is not as straightforward as with the independent set polytope since the order of arguments of ϕ matters, and the same vertex may be repeated among the arguments. For an ordered k-tuple of vertices $W = (w_1, \ldots, w_k) \in V(G)^k$ we thus define its characteristic vector χ^W of length $k|V(G)|$ by

$$\chi^W_{v,i} = \begin{cases} 1 & \text{if } v = w_i, \\ 0 & \text{otherwise.} \end{cases}$$

Note that χ^W always satisfies $\sum_{v \in V(G)} \chi^W_{v,i} = 1$ for each $i = 1, \ldots, k$, by the definition.

If $W = (w_1, \ldots, w_k) \in V(G)^k$ is such that $\phi(w_1, \ldots, w_k)$ holds true in G, we write $G \models \phi(w_1, \ldots, w_k)$. We can now give the following definition:

Definition 5.1 (FO polytope)**.** Let $\phi(x_1, \ldots, x_k)$ be an FO formula with k free variables. The *(first-order) ϕ-polytope of G*, denoted by $\mathrm{FOP}_\phi(G)$, is defined to be the convex hull of the characteristic vectors of every k-tuple of vertices of G such that $\phi(w_1, \ldots, w_k)$ holds true in G. That is,

$$\mathrm{FOP}_\phi(G) = \mathrm{conv}\left(\left\{\chi^W \in \{0,1\}^n \middle| \, W = (w_1, \ldots, w_k) \in V(G)^k,\ G \models \phi(w_1, \ldots, w_k)\right\}\right).$$

This definition is closely related to Definition 2.4 via the following observation:

Lemma 5.2. *Let $\iota(x_1, \ldots, x_k) \equiv \bigwedge_{i \neq j} \left(\neg edge(x_i, x_j) \wedge x_i \neq x_j\right)$ (the k-independent set formula). For every graph G, the ι-polytope $\mathrm{FOP}_\iota(G)$ is an extension of $\mathrm{STAB}_k(G)$.*

Proof. If G has n vertices then $\mathrm{STAB}_k(G) = \left\{ y \in \mathbb{R}^n \,\middle|\, y_v = \sum_{i=1}^k \chi^W_{v,i},\ \chi^W \in \mathrm{FOP}_\iota(G) \right\}$. Therefore, $\mathrm{STAB}_k(G)$ is a projection of $\mathrm{FOP}_\iota(G)$ given by the projection map described by $y_v = \sum_{i=1}^k \chi^W_{v,i}$ for all vertices v of G. □

Using the decomposition provided by Theorem 4.2, we can in fact obtain a more general result which we will discuss now. We will use the following weaker form[1] of a recent result of Kolman et al. [18].

Theorem 5.3 ((Kolman, Koutecký and Tiwary [18]))**.** *Let $\phi(x_1, \ldots, x_k)$ be an existential FO formula with k free variables and ℓ quantifiers. Then there exists a computable function $g : \mathbb{N} \times \mathbb{N} \to \mathbb{N}$, such that*

$$\mathrm{xc}\big(\mathrm{FOP}_\phi(G)\big) \leqslant g(k+\ell, \tau) \cdot n$$

holds for every integer n and every n-vertex graph G of treewidth τ. Furthermore, this extension can be computed in linear time for fixed k, ℓ and τ.

We are now able to directly extend Theorem 4.3 to the following restrictive fragment of FO logic. We say that an FO formula $\phi(x_1, \ldots, x_k)$ is *existential FO* if it can be written as $\phi(x_1, \ldots, x_k) \equiv \exists y_1 \ldots y_\ell \, \psi(x_1, \ldots, x_k, y_1, \ldots, y_\ell)$, where ψ is quantifier-free. The number ℓ of quantified variables in ϕ is called the *quantifier rank of ϕ*.

Theorem 5.4. *Let $\phi(x_1, \ldots, x_k)$ be an existential FO formula with k free variables and quantifier rank ℓ. Also, let $\mathcal{G}$ be any graph class of bounded expansion. Then there exists a computable function $f : \mathbb{N} \to \mathbb{N}$, depending on the expansion function of $\mathcal{G}$, such that*

$$\mathrm{xc}\big(\mathrm{FOP}_\phi(G)\big) \leqslant f(k+\ell) \cdot n$$

holds for every integer n and every n-vertex graph $G \in \mathcal{G}$. Furthermore, an explicit extension of $\mathrm{FOP}_\phi(G)$ of size at most $f(k+\ell) \cdot n$ can be found in linear time for fixed k, ℓ and $\mathcal{G}$.

Since the proof of Theorem 5.4 is a straightforward extension of the proof of Theorem 4.3, we skip it in this restricted paper.

Proof. We start with two simple facts of model theory:

a) If H is an induced subgraph of G, and $H \models \phi(w_1, \ldots, w_k)$ for $w_1, \ldots, w_k \in V(H)$, then $G \models \phi(w_1, \ldots, w_k)$ (since ϕ is existential).

b) If $G \models \phi(w_1, \ldots, w_k)$ for any $W = \{w_1, \ldots, w_k\} \subseteq V(G)$, then there is $U \subseteq V(G)$, $|U| \leq \ell$, such that $G[W \cup U] \models \phi(w_1, \ldots, w_k)$ where $G[W \cup U]$ is the subgraph of G induced on $W \cup U$ (since the quantifier rank of existential ϕ is $\leq \ell$).

[1] The result of Kolman et al. applies to Monadic Second Order logic: a logic that subsumes FO logic.

We can hence apply the same technique as in the proof of Theorem 4.3 – using a low treedepth coloring c of G now by $N_{\mathcal{G}}(k+\ell)$ colors from Theorem 4.2. Again, let $\mathcal{J}_{k+\ell} := \binom{[N_{\mathcal{G}}(k+\ell)]}{k+\ell}$ denote the set of $(k+\ell)$-element subsets of $[N_{\mathcal{G}}(k+\ell)]$, and let a subgraph $G_J \subseteq G$ where $J \in \mathcal{J}_{k+\ell}$, be defined as the subgraph of G induced on $\bigcup_{j\in\mathcal{J}_{k+\ell}} c^{-1}(j)$ – the color classes of c indexed by J. By a),b) we immediately get

$$\mathrm{FOP}_\phi(G) = \mathrm{conv}\left(\bigcup\nolimits_{J\in\mathcal{J}_{k+\ell}} \mathrm{FOP}_\phi(G_J)\right).$$

From Theorems 2.2 and 5.3 (via Observation 4.1) we then analogously conclude

$$\begin{aligned}\mathrm{xc}\big(\mathrm{FOP}_\phi(G)\big) &\leqslant |\mathcal{J}_{k+\ell}| + \sum_{J\in\mathcal{J}_{k+\ell}} \mathrm{xc}\big(\mathrm{FOP}_\phi(G_J)\big)\\ &\leqslant |\mathcal{J}_{k+\ell}| \cdot \big(1 + g(k+\ell, k+\ell-1)\cdot n\big) \leqslant f(k+\ell)\cdot n\,.\end{aligned}$$

Again, this extended formulation can be constructed in linear time for fixed k, ℓ. □

6. Conclusions

We have begun to study the question: to which extent FPT tractability of the k-independent set problem on graph classes is related to the FPT extension complexity of the (corresponding) k-independent-set polytope? Not surprisingly, we confirm that there cannot be FPT extensions of this polytope in the class of all graphs (note, though, that our proof is absolute and does not rely on the assumption $FPT \neq W[1]$). On the other hand, the k-independent-set problem is linear-time FPT tractable on graph classes of bounded expansion [13], and we construct a linear FPT extension for its polytope on such classes. This positive result easily carries over to all existential FO problems of graphs.

There are two possible directions of future research in this regard:

1. A graph class $\mathcal{G}$ is *nowhere dense* [16] if there is no integer d such that $\mathcal{G}\nabla d$ contains all graphs. Every graph class of bounded expansion is nowhere dense, but the converse is not true. The k-independent-set problem is also known to be in FPT on every nowhere dense class, which follows from a more general result of [19]. It is natural to ask whether the same can hold for the fixed-parameter extension complexity of its polytope.

2. The deep tractability results of [20, 19] (on bounded expansion and nowhere dense classes, respectively) address problems in full FO logic of graphs, i.e., allowing also for universal quantifiers in the problem expression. This suggests that perhaps, for every FO formula ϕ, the related ϕ-polytope can also have FPT extension complexity on, say, graph classes of bounded expansion. Though, the involved proof techniques of [20, 19] do not seem to easily translate to the extension complexity setting.

In a broader view, one may regard the property of (the problem polytope) having an FPT extension complexity as a finer (case-by-case) resolution of the class FPT. For an explanation; the well-established assumption $FPT \neq W[1]$ implies that problems not in FPT do not have FPT extensions, while on the other hand the example of the matching polytope [7] suggests that there may also be FPT problems whose polytopes do not have FPT extensions. We believe this task is worth further detailed investigation.

One may try to proceed even further and ask a general question:

3. Is it true that all $W[t]$-hard problems for some $t \geq 1$ do not admit FPT extensions?

However, this question is not even easy to formulate since the polytope we associate with a problem remains a specific choice which by no means is the only choice.

Moreover, it can be argued that either possible answer to the very broad question (3) would be a significant breakthrough in the complexity world. Say, since some FPT problems such as k-vertex cover do admit FPT extensions [12], an affirmative answer to (3) would imply that this problem is

not $W[t]$-complete and so $FPT \neq W[t]$. On the other hand, if the answer to (3) was no, then this would imply the existence of non-uniform FPT circuits for $W[t]$-complete problems which is considered unlikely.

References

[1] M. Conforti, G. Cornuéjols, G. Zambelli, Extended formulations in combinatorial optimization, 4OR 8 (2010) 1–48.

[2] F. Vanderbeck, L. A. Wolsey, Reformulation and decomposition of integer programs, in: M. J. et al. (Ed.), 50 Years of Integer Programming 1958-2008, Springer, 2010, pp. 431–502.

[3] V. Kaibel, Extended formulations in combinatorial optimization, Optima 85 (2011) 2–7.

[4] L. A. Wolsey, Using extended formulations in practice, Optima 85 (2011) 7–9.

[5] S. Fiorini, S. Massar, S. Pokutta, H. R. Tiwary, R. de Wolf, Exponential lower bounds for polytopes in combinatorial optimization, J. ACM 62 (2) (2015) 17.

[6] G. Braun, S. Fiorini, S. Pokutta, D. Steurer, Approximation limits of linear programs (beyond hierarchies), Math. Oper. Res. 40 (3) (2015) 756–772. doi:10.1287/moor.2014.0694.

[7] T. Rothvoß, The matching polytope has exponential extension complexity, in: Proc. of the 46th ACM Symposium on Theory of Computing, (STOC), 2014, pp. 263–272.

[8] S. Chan, J. Lee, P. Raghavendra, D. Steurer, Approximate constraint satisfaction requires large LP relaxations, FOCS'13 (2013).

[9] D. Avis, H. R. Tiwary, On the extension complexity of combinatorial polytopes, in: Proc. ICALP(1) 2013, 2013, pp. 57–68.

[10] S. Pokutta, M. Van Vyve, A note on the extension complexity of the knapsack polytope, Operations Research Letters 41 (4) (2013) 347–350.

[11] R. G. Downey, M. R. Fellows, Fundamentals of Parameterized Complexity, Texts in Computer Science, Springer, 2013. doi:10.1007/978-1-4471-5559-1.

[12] A. Buchanan, Extended formulations for vertex cover, available on Optimization Online (2015).

[13] J. Nešetřil, P. Ossona de Mendez, Linear time low tree-width partitions and algorithmic consequences, in: STOC'06, ACM, 2006, pp. 391–400. doi:10.1145/1132516.1132575.

[14] E. Balas, Disjunctive programming: Properties of the convex hull of feasible points, Discrete Applied Mathematics 89 (1–3) (1998) 3–44.

[15] A. Buchanan, S. Butenko, Tight extended formulations for independent set, available on Optimization Online (2014).

[16] J. Nešetřil, P. Ossona de Mendez, Sparsity: Graphs, Structures, and Algorithms, Vol. 28 of Algorithms and Combinatorics, Springer, 2012.

[17] J. Nešetřil, P. Ossona de Mendez, Grad and classes with bounded expansion I. Decompositions, European J. Combin. 29 (3) (2008) 760–776.

[18] P. Kolman, M. Koutecký, H. R. Tiwary, Extension complexity, mso logic, and treewidth, CoRR abs/1507.04907.
URL http://arxiv.org/abs/1507.04907

[19] M. Grohe, S. Kreutzer, S. Siebertz, Deciding first-order properties of nowhere dense graphs, in: STOC'14, ACM, 2014, pp. 89–98. doi:10.1145/2591796.2591851.

[20] Z. Dvořák, D. Král', R. Thomas, Deciding first-order properties for sparse graphs, in: FOCS'10, IEEE Computer Society, 2010, pp. 133–142.